A Linear-time Algorithm for Sparsification of Unweighted Graphs

Ramesh Hariharan
Strand Life Sciences

Debmalya Panigrahi
CSAIL, MIT

Abstract

Given an undirected graph G and an error parameter $\varepsilon > 0$, the *graph sparsification* problem requires sampling edges in G and giving the sampled edges appropriate weights to obtain a sparse graph G_ε with the following property: the weight of every cut in G_ε is within a factor of $(1 \pm \varepsilon)$ of the weight of the corresponding cut in G. If G is unweighted, an $O(m \log n)$-time algorithm for constructing G_ε with $O(n \log n/\varepsilon^2)$ edges in expectation, and an $O(m)$-time algorithm for constructing G_ε with $O(n \log^2 n/\varepsilon^2)$ edges in expectation have recently been developed [9]. In this paper, we improve these results by giving an $O(m)$-time algorithm for constructing G_ε with $O(n \log n/\varepsilon^2)$ edges in expectation, for unweighted graphs. Our algorithm is optimal in terms of its time complexity; further, no efficient algorithm is known for constructing a sparser G_ε. Our algorithm is Monte-Carlo, i.e. it produces the correct output with high probability, as are all efficient graph sparsification algorithms.

1 Introduction

A *cut* of an undirected graph is a partition of its vertices into two disjoint sets. The *weight* of a cut is the sum of weights of the edges crossing the cut, i.e. edges having one endpoint each in the two vertex subsets of the partition. For unweighted graphs, each edge is assumed to have unit weight. Cuts play an important role in many problems in graphs: e.g., the maximum flow between a pair of vertices is equal to the minimum weight cut separating them.

A *skeleton* G' of an undirected graph G is a subgraph of G on the same set of vertices where each edge in G' can have an arbitrary weight. The problem of finding an appropriately weighted sparse skeleton for an undirected graph G that approximately preserves the weights of all cuts in G was introduced and studied by Karger *et al* in a series of results [11, 12, 3] culminating in the following theorem. Throughout this paper, for any undirected graph G and any $\varepsilon \in (0,1]$, $(1 \pm \varepsilon)G$ will denote the set of all appropriately weighted subgraphs of G where the weight of every cut in the subgraph is within a factor of $(1 \pm \varepsilon)$ of the weight of the corresponding cut in G.

Theorem 1 (Benczúr-Karger [3])**.** *For any undirected graph G with m edges and n vertices, and for any error parameter $\varepsilon \in (0,1]$, a skeleton G_ε containing $O(\frac{n \log n}{\varepsilon^2})$ edges in expectation such that $G_\varepsilon \in (1 \pm \varepsilon)G$ with high probability*[1] *can be found in $O(m \log^2 n)$ time if G is unweighted and $O(m \log^3 n)$ time otherwise.*

Besides its combinatorial ramifications, the importance of this result stems from its use as a pre-processing step in several graph algorithms, e.g. to obtain an $\tilde{O}(n^{3/2} + m)$-time algorithm for approximate maximum flow using the $\tilde{O}(m^{3/2})$-time algorithm for exact maxflow due to Goldberg and Rao [7]; and more recently, $\tilde{O}(n^{3/2} + m)$-time algorithms for approximate sparsest cut [13, 19].

Subsequent to Benczúr and Karger's work, Spielman and Teng [21, 22] extended their results to preserving all quadratic forms, of which cuts are a special case; however, the size of the skeleton constructed was $O(n \log^c n)$ for some large constant c. Spielman and Srivastava [20] improved this result by constructing skeletons of size $O(\frac{n \log n}{\varepsilon^2})$ in $O(m \log^{O(1)} n)$ time, while continuing to preserve all quadratic forms. Recently, this result was further improved by Batson *et al* [2] who gave a deterministic algorithm for constructing skeletons of size $O(\frac{n}{\varepsilon^2})$ that preserve the weights of all cuts whp. While their result is optimal in terms of the size of the skeleton constructed, the time complexity of their algorithm is $O(\frac{mn^3}{\varepsilon^2})$, rendering it somewhat useless in terms of applications.

Recently, further progress has been made on efficiently constructing a skeleton graph in the form of the following theorem due to Hariharan and Panigrahi [9].

Theorem 2 (Hariharan-Panigrahi [9])**.** *For an undirected graph G with m edges and n vertices, and for any error parameter $\varepsilon \in (0,1]$, the following algorithmic results can be obtained for constructing a skeleton graph G_ε that is in $(1 \pm \varepsilon)G$ with high probability:*

- *If the expected number of edges in G_ε is $O(n \log^2 n / \varepsilon^2)$, then G_ε can be constructed in $O(m)$ time if G has polynomial edge weights and $O(m \log^2 n)$ time if G has arbitrary edge weights.*
- *If the expected number of edges in G_ε is $O(n \log n / \varepsilon^2)$, then G_ε can be constructed in $O(m \log n)$ time if G is unweighted, and $O(m \log^2 n)$ time if G has polynomial edge weights.*

[1] We say that a property holds *with high probability* (or whp) for a graph on n vertices if its failure probability can be bounded by the inverse of a fixed polynomial in n.

Combining the above two results, one can obtain an algorithm to construct a skeleton graph G_ε that preserves the weights of all cuts whp and has $O(n\log n/\varepsilon^2)$ edges in expectation in $O(m + n\log^4 n/\varepsilon^2)$ time if G has polynomial edge weights.

A natural conclusion for this line of work would be to obtain an $O(m)$-time algorithm for constructing a skeleton graph G_ε containing $O(n\log n/\varepsilon^2)$ edges in expectation. In this correspondence, we obtain this result in the form of the following theorem if G is unweighted. It may be noted that even if G is unweighted, the best sparsification result known previously was Theorem 2.

Theorem 3. *For an undirected unweighted graph G with m edges and n vertices, and for any error parameter $\varepsilon \in (0,1]$, a skeleton graph G_ε that is in $(1\pm\varepsilon)G$ with high probability and has $O(n\log^2 n/\varepsilon^2)$ edges in expectation, can be constructed in $O(m)$ time.*

Note that the above algorithm is optimal in terms of its running time, and there is no efficient algorithm known for constructing a sparser skeleton, even for unweighted graphs. (As mentioned previously, the only algorithm known that constructs a sparser skeleton has a time complexity of $O(n^3 m/\varepsilon^2)$ [2].)

Before describing our algorithm in more detail, it is worth mentioning some of the related research in graph sparsification. In a recent result, Fung and Harvey [5] show that sampling uniformly random spanning trees of a graph produces good sparsifiers. This approach was used previously to obtain coarser sparsifiers by Goyal *et al* [8]. Fung and Harvey also show that sampling edges according to their standard connectivities also produces good sparsifiers, a result obtained independently by Hariharan and Panigrahi [9]. The problem of graph sparsification in the semi-streaming model was first considered by Ahn and Guha [1] who gave a one-pass algorithm for constructing a skeleton G_ε containing $O(n\log n\log(m/n)/\varepsilon^2)$ edges. Recently, Goel *et al* have given the following algorithms for this problem [6]:

- An $O(m\log\log n)$-time one-pass algorithm for constructing a skeleton graph G_ε containing $O(n\log^2 n/\varepsilon^2)$ edges in expectation. The size of the skeleton can be improved to $O(n\log n/\varepsilon^2)$ edges; however, the time complexity of the algorithm then becomes $O(m\log\log n + n\log^5 n/\varepsilon^2)$.

- An $O(m)$-time two-pass algorithm for constructing a skeleton graph G_ε containing $O(n\log n/\varepsilon^2)$ edges in expectation, if $m = \Omega(n^{1+\delta})$ for some constant $\delta > 0$.

Observe that both results, if applied to a non-streaming model, are weaker than Theorem 2. Another area of recent interest, though not directly related to our problem, is that of vertex sparsification [15, 14]. Given a graph $G = (V,E)$ and a subset of vertices $S \subset V$, the goal here is to create a graph $G_S = (S,E_S)$ that approximately preserves some desired connectivity property of G (e.g. minimum steiner cut [15], maximum multi-commodity flow [14]).

1.1 Our Techniques

All previous algorithms for graph sparsification have two phases: in the first phase, a suitable probability p_e for sampling each edge e is determined; and then, in the second phase, every edge is independently sampled with probability p_e and given weight $1/p_e$ in the skeleton graph if selected in the sample. Our main technical novelty is in interleaving the sampling process with that of estimating sampling probabilities. Such interleaving leads to several technical hurdles:

- It introduces dependence between the sampling of different edges. Such dependence has appeared previously in sparsification algorithms for the semi-streaming model, but the nature of the dependence in our algorithm is somewhat different from that in the streaming algorithms.

- An edge may now be sampled multiple times, and errors are accrued in each such sampling. This requires us to choose the interleaved sampling probabilities very carefully so that the errors do not add up.

At a high level, our algorithm has the same iterative structure as algorithms in [3] and [9]. In each iteration, the algorithm identifies suitable sampling probabilities of a subset of edges and removes them from the graph. It is in what the algorithm does with the remaining edges that our algorithm differs from previous work. While all remaining edges are retained for the next iteration in previous algorithms, we sample all the edges with probability 1/2 and retain only half of them in expectation for the next iteration. The intuition behind this sampling comes from the observation that the sampling probabilities decrease (approximately by a factor of 2) with every iteration; therefore, a natural approach is to sample the remaining edges with probability 1/2 and retain only the selected edges thereby reducing the time complexity of the next iteration.

Suppose X_i be the set of remaining edges at the beginning of iteration i, F_i be the set of edges whose sampling probabilities are determined in iteration i and $Y_i = X_i \setminus F_i$ be the set of remaining edges after iteration i. (Note that X_{i+1} is therefore constructed by sampling each edge in Y_i with probability 1/2.) Our proof technique consists of two parts. First, we show that the graph S containing appropriately weighted edges in $\cup_i F_i$ is in $(1 \pm \varepsilon/3)G$ whp, i.e. even though edges in $Y_i \setminus X_{i+1}$ are sampled out between iterations i and $i+1$ for each i, the retained edges (when weighted appropriately) are sufficient to preserve all cuts. In the second step of the proof, we show that the skeleton graph G_ε constructed by sampling edges in $\cup_i F_i$ and giving them appropriate weights is in $(1 \pm \varepsilon/3)S$ whp. For this proof, we use the generic sparsification framework developed recently by Hariharan and Panigrahi [9]. Combining these two steps, we conclude that $G_\varepsilon \in (1 \pm \varepsilon)G$ whp.

Roadmap. In section 2, we give an outline of the generic sampling framework from [9] that we use later in our proof. In section 3, we describe our sparsification algorithm, prove its correctness and derive its time complexity. Finally, we conclude with some open questions in section 4.

2 Preliminaries

We first need to introduce the notion of k-heavy edges, for any $k > 0$.

Definition 1. *An edge $e = (u, v)$ of an undirected graph $G = (V, E)$ is said to be k-heavy if the maximum flow between vertices u and v in G is at least k.*

By Menger's theorem (see e.g., [4]), it follows that if $e = (u, v)$ is k-heavy, then the weight of every cut in G having u and v on different sides is at least k.

2.1 Outline of Sparsification Framework from [9]

Suppose $G = (V, E)$ is an undirected graph where edge $e \in E$ has a positive integer weight w_e. Let $G_M = (V, E_M)$ denote the multi-graph constructed by replacing each edge e by w_e unweighted parallel edges $e_1, e_2, \ldots, e_{w_e}$. Consider any $\varepsilon \in (0, 1]$. Suppose we construct a skeleton G_ε where each edge $e_\ell \in E_M$ is present in graph G_ε independently with probability p_e, and if present, it is given a weight of $1/p_e$. Let $p_e = \min(\frac{96\alpha \ln n}{0.38\lambda_e \varepsilon^2}, 1)$, where α is independent of e and λ_e is some parameter of e satisfying $\lambda_e \leq 2^n - 1$. The authors describe a sufficient condition that characterizes a good choice of α and λ_e's.

To describe this sufficient condition, partition the edges in G_M according to the value of λ_e into sets $R_0, R_1, \ldots, R_K$ where $K = \lfloor \lg \max_{e \in E}\{\lambda_e\} \rfloor \leq n - 1$ and $e_i \in R_j$ iff $2^j \leq \lambda_e \leq 2^{j+1} - 1$. Now, let $\mathbf{Q} =$

$(Q_0, Q_1, Q_2, \ldots, Q_i = (V, W_i), \ldots, Q_k)$ be a sequence of subgraphs of G_M (edges of G_M are allowed to be replicated multiple times in the Q_is) such that $R_i \subseteq W_i$ for every i. **Q** is said to be a (π, α)-*certificate* corresponding to the above choice of α and λ_e's if the following properties are satisfied:

π-connectivity For $i \geq 0$, any edge $e_\ell \in R_i$ is π-heavy in Q_i.

α-overlap For any cut C containing c edges in G_M, let $w_i^{(C)}$ be the number of edges that cross C in Q_i. Then, for all cuts C, $\sum_{i=0}^{k} \frac{w_i^{(C)} 2^{i-1}}{\pi} \leq \alpha c$.

Then, the following theorem holds.

Theorem 4 (Hariharan-Panigrahi [9] (Theorem 8))**.** *If there exists a (π, α)-certificate for a particular choice of α and λ_e's , then the skeleton $G_\varepsilon \in (1 \pm \varepsilon)G$ with probability at least $1 - 4/n$. Further G_ε has $O(\frac{\alpha \log n}{\varepsilon^2} \sum_{e \in E} \frac{w_e}{\lambda_e})$ edges in expectation.*

We also need the following lemma, which is a slight variation of Lemma 5 from [9]. (For completeness, we give a proof in the appendix.) For an undirected unweighted graph $G = (V, E)$, let $R \subseteq E$ and $Q \supseteq R$ be subsets of edges such that R is π-heavy in (V, Q). Suppose each edge $e \in R$ is sampled with probability p, and if selected, given a weight of $1/p$ to form a set of weighted edges $\widehat{R}$. Now, for any cut C in G, let $R^{(C)}$, $Q^{(C)}$ and $\widehat{R^{(C)}}$ be the sets of edges crossing cut C in R, Q and $\widehat{R}$ respectively; also let the total weight of edges in $R^{(C)}$, $Q^{(C)}$ and $\widehat{R^{(C)}}$ be $r^{(C)}$, $q^{(C)}$ and $\widehat{r^{(C)}}$ respectively. Then the following lemma holds.

Lemma 1. *For any $\delta \in (0, 1]$ satisfying $\delta^2 \cdot p \cdot \pi \geq \frac{6 \ln n}{0.38}$,*

$$|r^{(C)} - \widehat{r^{(C)}}| \leq \delta q^{(C)}$$

for all cuts C in G with probability at least $1 - 4/n^2$.

2.2 Nagamochi-Ibaraki Forests

We first introduce the notion of *spanning forests* of a graph. As earlier, G denotes a graph with integer edge weights w_e for edge e and G_M is the unweighted multi-graph where e is replaced with w_e parallel unweighted edges.

Definition 2. *A* spanning forest *T of G_M (or equivalently of G) is an (unweighted) acyclic subgraph of G satisfying the property that any two vertices are connected in T if and only if they are connected in G.*

We partition the set of edges in G_M into a set of forests $T_1, T_2, \ldots$ using the following rule: *T_i is a spanning forest of the graph formed by removing all edges in $T_1, T_2, \ldots, T_{i-1}$ from G_M such that for any edge $e \in G$, all its copies in G_M appear in a set of contiguous forests $T_{i_e}, T_{i_e+1}, \ldots, T_{i_e+w_e-1}$.* This partitioning technique was introduced by Nagamochi and Ibaraki in [18], and these forests are known as *Nagamochi-Ibaraki forests* (or NI forests). The following is a basic property of NI forests.

Lemma 2 (Nagamochi-Ibaraki [18, 17])**.** *For any pair of vertices u, v, they are connected in NI forests $T_1, T_2, \ldots, T_{k(u,v)}$ for some $k(u, v)$ and not connected in any forest T_j, for $j > k(u, v)$.*

Nagamochi and Ibaraki also gave an algorithm for constructing NI forests that runs in $O(m + n)$ time if G_M is a simple graph (i.e. G is unweighted) and $O(m + n \log n)$ time otherwise [18, 17].

3 The Algorithm

We describe out sparsification algorithm for an unweighted graph $G = (V,E)$ with m edges and n vertices, and an error parameter $\varepsilon \in (0,1]$ as inputs. We prove that the skeleton graph G_ε produced by the algorithm is in $(1 \pm \varepsilon)G$ with high probability. We then show that the expected number of edges in G_ε is $O(n \log n/\varepsilon^2)$. Finally, we prove that the expected time complexity of the algorithm is $O(m)$.

Description of the Algorithm. The algorithm has three phases. The first phase has the following steps:

- If $m \leq 2\rho n$, where $\rho = \frac{1014 \ln n}{0.38\varepsilon^2}$, G is sparse enough itself. Therefore, we take G as our skeleton graph.
- Otherwise, we construct a set of NI forests of G and all edges in the first 2ρ NI forests are included in the skeleton graph G_ε with weight 1. We call these edges F_0. The edge set Y_0 is then defined as $E \setminus F_0$.

The second phase is iterative. The input to iteration i is a graph (V, Y_{i-1}), which is a subgraph of the input graph to iteration $i-1$ (i.e. $Y_{i-1} \subseteq Y_{i-2}$). Iteration i comprises the following steps:

- If the number of edges in Y_{i-1} is at most $2\rho n$, we take all those edges in G_ε with weight 2^{i-1} each, and terminate the algorithm.
- Otherwise, all edges in Y_i are sampled with probability $1/2$; call the sample X_i and let $G_i = (V, X_i)$.
- We identify a set of edges in X_i (call this set F_i) that has the following properties:
  - The number of edges in F_i is at most $2k_i|V_c|$, where $k_i = \rho \cdot 2^{i+1}$, and V_c is the set of components in (V, Y_i), where $Y_i = X_i \setminus F_i$.
  - Each edge in Y_i is k_i-heavy in G_i.
- We give a sampling probability $p_i = \min(\frac{3}{169 \cdot 2^{2i-9}}, 1)$ to all edges in F_i.

The final phase consists of replacing each edge in F_i (for each i) with 2^i parallel edges, and then sampling each parallel edge independently with probability p_i. If an edge is selected in the sample, it is added to the skeleton graph G_ε with weight $1/p_i$.

We now give a short description of the sub-routine that constructs the set F_i in iteration i of the second phase of the algorithm. This sub-routine is iterative itself: we start with $V_c = V$ and $E_c = X_i$, and let $G_c = (V_c, E_c)$. We repeatedly construct $k_i + 1$ NI forests for G_c where $k_i = \rho 2^{i+1} + 1$ and contract all edges in the $(k_i + 1)$st forest to obtain a new G_c, until $|E_c| \leq \rho 2k_i|V_c|$. The set of edges E_c that finally achieves this property forms F_i.

The complete algorithm is given in Figure 1.

Cut Preservation. We first show that the skeleton graph G_ε produced by the above algorithm is in $(1 \pm \varepsilon)G$ with high probability. We use the following notation throughout: for any set of unweighted edges Z, cZ denotes these edges with a weight of c given to each edge.

Our goal is to prove the following theorem.

Theorem 5. *$G_\varepsilon \in (1 \pm \varepsilon)G$ with probability at least $1 - 8/n$.*

As outlined in the introduction, our proof has two stages. Let K be the maximum value of i for which $F_i \neq \emptyset$; let $S = \left(\cup_{i=0}^{K} 2^i F_i\right) \cup 2^K Y_K$ and $G_S = (V, S)$. Then, in the first stage, we prove the following theorem.

1. Set $\rho = \frac{1014 \ln n}{0.38\varepsilon^2}$.
2. If $m \leq 2\rho n$, then $G_\varepsilon = G$; else, go to step 3.
3. Construct NI forests $T_1, T_2, \ldots$ for G.
4. Set $i = 0$.
5. Set $X_i = E$; $F_i = \cup_{1 \leq j \leq 2\rho} T_j$; $Y_i = X_i \setminus F_i$.
6. Add each edge in F_i to G_ε with weight 1.
7. If $|Y_i| \leq 2\rho n$, then add each edge in Y_i to G_ε with weight 2^{i-1} and terminate; else, go to step 8.
8. Sample each edge in Y_i with probability 1/2 to construct X_{i+1}.
9. Increment i by 1.
10. Set $E_c = X_i$; $V_c = V$.
11. Set $k_i = \rho \cdot 2^{i+1}$.
12. If $|E_c| \leq 2k_i|V_c|$, then
    (a) Set $F_i = E_c$; $Y_i = X_i \setminus E_c$.
    (b) For each edge $e \in F_i$, set $\lambda_e = \rho \cdot 4^i$.
    (c) Go to step 7.

    Else,

    (a) Construct NI forests $T_1, T_2, \ldots, T_{k_i+1}$ for graph $G_c = (V_c, E_c)$.
    (b) Update $G_c = (V_c, E_c)$ by contracting all edges in T_{k_i+1}.
    (c) Go to step 12.
13. For each edge $e \in \cup_i F_i$,
    (a) Set $p_e = \min(\frac{9216 \ln n}{0.38\lambda_e \varepsilon^2}, 1) = \min(\frac{3}{169 \cdot 2^{2i-9}}, 1)$.
    (b) Generate r_e from **Binomial**$(2^i, p_e)$.
    (c) If $r_e > 0$, add edge e to G_ε with weight r_e/p_e.

Figure 1: Our sparsification algorithm

Theorem 6. *$G_S \in (1 \pm \varepsilon/3)G$ with probability at least $1 - 4/n$.*

In the second stage, we prove the following theorem.

Theorem 7. *$G_\varepsilon \in (1 \pm \varepsilon/3)G_S$ with probability at least $1 - 4/n$.*

Combining the above two theorems and using the union bound, we obtain Theorem 5. (Observe that since $\varepsilon \leq 1$, $(1+\varepsilon/3)^2 \leq 1+\varepsilon$ and $(1-\varepsilon/3)^2 \geq 1-\varepsilon$).

The following property is key to proving both Theorem 6 and Theorem 7.

Lemma 3. *For any $i \geq 0$, any edge $e \in Y_i$ is k_i-heavy in $G_i = (V, X_i)$, where $k_i = \rho \cdot 2^{i+1}$.*

Proof. For $i = 0$, all edges in Y_i are in NI forests $T_{2\rho+1}, T_{2\rho+2}, \ldots$ of $G_i = G$. The proof follows from Lemma 2.

We now prove the lemma for $i \geq 1$. Let $G_e = (V_e, E_e)$ be the component of G_i containing e. We will show that e is k_i-heavy in G_e; since G_e is a subgraph of G_i, the lemma follows. In the execution of the else block of step 12 on G_e, there are multiple contraction operations, each of them comprising the contraction of a set of edges. We show that any such contracted edge is k_i-heavy in G_e; it follows that e is k_i-heavy in G_e.

Let G_e have t contraction phases and let the graph produced after contraction phase r be $G_{e,r}$. We now prove that all edges contracted in phase r must be k_i-heavy in G_e by induction on r. For $r = 1$, since e appears in the (k_i+1)st NI forest of phase 1, e is k_i-heavy in G_e by Lemma 2. For the inductive step, assume that the property holds for phases $1, 2, \ldots, r$. Any edge that is contracted in phase $r+1$ appears in the (k_i+1)st NI forest of phase $r+1$; therefore, e is k_i-connected in $G_{e,r}$ by Lemma 2. By the inductive hypothesis, all edges of G_e contracted in previous phases are k_i-heavy in G_e; therefore, an edge that is k_i-heavy in $G_{e,r}$ must have been k_i-heavy in G_e. $\square$

Proof of Theorem 6. The next lemma follows from Lemma 1.

Lemma 4. *With probability at least $1 - 4/n^2$, for every cut C in G_i, $|2x_{i+1}^{(C)} + f_i^{(C)} - x_i^{(C)}| \leq \frac{\varepsilon/13}{2^{i/2}} \cdot x_i^{(C)}$.*

Proof. Use the following parameters in Lemma 1:

- $R = Y_i$; $Q = X_i$; $\widehat{R} = 2X_{i+1}$
- $\delta = \frac{\varepsilon/13}{2^{i/2}}$; $p = 1/2$; $\pi = \rho \cdot 2^{i+1}$.

Lemma 3 ensures that R is π-heavy in (V, Q); also, it can be verified that $\delta^2 \cdot p \cdot \pi = 6 \ln n$. $\square$

We use the above lemma to prove the following lemma.

Lemma 5. *Let $S_j = \left(\cup_{i=j}^{K} 2^{i-j} F_i\right) \cup 2^{K-j} Y_K$ for any $j \geq 0$. Then, $S_j \in (1 \pm (\varepsilon/3)2^{-j/2})G_j$ with probability at least $1 - 4/n$, where $G_j = (V, X_j)$.*

To prove this lemma, we need to use the following fact.

Fact 1. *Let $x \in (0, 1]$ and $r_i = 13 \cdot 2^{i/2}$. Then, for any $k \geq 0$,*

$$\begin{aligned} \prod_{i=0}^{k} (1 + x/r_i) &\leq 1 + x/3 \\ \prod_{i=0}^{k} (1 - x/r_i) &\geq 1 - x/3. \end{aligned}$$

Proof. We prove by induction on k. For $k=0$, the property trivially holds. Suppose the property holds for $k-1$. Then,

$$\begin{aligned}
\prod_{i=0}^{k}(1+x/r_i) &= \prod_{i=0}^{k}(1+\frac{x}{13\cdot 2^{i/2}}) \\
&= (1+x/13)\cdot\prod_{i=1}^{k}\left(1+\frac{x/\sqrt{2}}{13\cdot 2^{(i-1)/2}}\right) \\
&\le (1+x/13)\cdot(1+x/(3\sqrt{2})) \\
&\le 1+x/3 \\
\prod_{i=0}^{k}(1-x/r_i) &= \prod_{i=0}^{k}(1-\frac{x}{13\cdot 2^{i/2}}) \\
&= (1-x/13)\cdot\prod_{i=1}^{k}\left(1-\frac{x/\sqrt{2}}{13\cdot 2^{(i-1)/2}}\right) \\
&\ge (1-x/13)\cdot(1-x/(3\sqrt{2})) \\
&\ge 1-x/3.
\end{aligned}$$

□

Proof of Lemma 5. For any cut C in G, let the edges crossing C in S_j be $S_j^{(C)}$, and let their total weight be $s_j^{(C)}$. Also, let $X_i^{(C)}$, $Y_i^{(C)}$ and $F_i^{(C)}$ be the set of edges crossing cut C in X_i, Y_i and F_i respectively, and let their total weights be $x_i^{(C)}$, $y_i^{(C)}$ and $f_i^{(C)}$. (Recall that all edges in X_i, Y_i and F_i are unweighted; therefore $x_i^{(C)}=|X_i^{(C)}|$, $y_i^{(C)}=|Y_i^{(C)}|$ and $f_i^{(C)}=|F_i^{(C)}|$.)

Since $K\le n-1$, we can use the union bound on Lemma 4 to conclude that with probability at least $1-4/n$, for every $0\le i\le K$ and for all cuts C,

$$\begin{aligned}
2x_{i+1}^{(C)}+f_i^{(C)} &\le (1+\varepsilon/r_i)x_i^{(C)} \\
2x_{i+1}^{(C)}+f_i^{(C)} &\ge (1-\varepsilon/r_i)x_i^{(C)},
\end{aligned}$$

where $r_i=13\cdot 2^{i/2}$. Then,

$$\begin{aligned}
s_j^C &= 2^{K-j}y_K^{(C)}+2^{K-j}f_K^{(C)}+2^{K-1-j}f_{K-1}^{(C)}+\ldots+f_j^{(C)} \\
&= 2^{K-j}x_K^{(C)}+2^{K-1-j}f_{K-1}^{(C)}+\ldots+f_j^{(C)} \quad \text{since } y_K^{(C)}+f_K^{(C)}=x_K^{(C)} \\
&= 2^{K-1-j}(2x_K^{(C)}+f_{K-1}^{(C)})+(2^{K-2-j}f_{K-2}^{(C)}+\ldots+f_j^{(C)}) \\
&\le (1+\varepsilon/r_{K-1})2^{K-1-j}x_{K-1}^{(C)}+(2^{K-2-j}f_{K-2}^{(C)}+\ldots+f_j^{(C)}) \\
&\le (1+\varepsilon/r_{K-1})(2^{K-1-j}x_{K-1}^{(C)}+2^{K-2-j}f_{K-2}^{(C)}+\ldots+f_j^{(C)}) \\
&\ldots \\
&\le (1+\varepsilon/r_{K-1})(1+\varepsilon/r_{K-2})\ldots(1+\varepsilon/r_j)x_j^{(C)} \\
&\le (1+(\varepsilon 2^{-j/2})/r_{K-1-j})(1+(\varepsilon 2^{-j/2})/r_{K-2-j})\ldots(1+(\varepsilon 2^{-j/2})/r_0)x_j^{(C)} \quad \text{since } r_{j+i}=r_i\cdot 2^{j/2} \\
&\le (1+(\varepsilon/3)2^{-j/2})x_j^{(C)} \quad \text{by Fact 1.}
\end{aligned}$$

Similarly,

$$
\begin{aligned}
s_j^C &= 2^{K-j} y_K^{(C)} + 2^{K-j} f_K^{(C)} + 2^{K-1-j} f_{K-1}^{(C)} + \ldots + f_j^{(C)} \\
&= 2^{K-j} x_K^{(C)} + 2^{K-1-j} f_{K-1}^{(C)} + \ldots + f_j^{(C)} \quad \text{since } y_K^{(C)} + f_K^{(C)} = x_K^{(C)} \\
&= 2^{K-1-j}(2x_K^{(C)} + f_{K-1}^{(C)}) + (2^{K-2-j} f_{K-2}^{(C)} + \ldots + f_j^{(C)}) \\
&\geq (1-\varepsilon/r_{K-1}) 2^{K-1-j} x_{K-1}^{(C)} + (2^{K-2-j} f_{K-2}^{(C)} + \ldots + f_j^{(C)}) \\
&\geq (1-\varepsilon/r_{K-1})(2^{K-1-j} x_{K-1}^{(C)} + 2^{K-2-j} f_{K-2}^{(C)} + \ldots + f_j^{(C)}) \\
&\ldots \\
&\geq (1-\varepsilon/r_{K-1})(1-\varepsilon/r_{K-2}) \ldots (1-\varepsilon/r_j) x_j^{(C)} \\
&\geq (1-(\varepsilon 2^{-j/2})/r_{K-1-j})(1-(\varepsilon 2^{-j/2})/r_{K-2-j}) \ldots (1-(\varepsilon 2^{-j/2})/r_0) x_j^{(C)} \quad \text{since } r_{j+i} = r_i \cdot 2^{j/2} \\
&\geq (1-(\varepsilon/3) 2^{-j/2}) x_j^{(C)} \quad \text{by Fact 1.}
\end{aligned}
$$

□

Theorem 6 now follows as a corollary of the above lemma for $j = 0$.

Proof of Theorem 7. Now, we use the sparsification framework developed in [9] and outlined previously in section 2 to prove Theorem 7. Observe that edges $F_0 \cup 2^K Y_K$ are identical in G_S and G_ε. Therefore, we do not consider these edges in the analysis below.

For any $i \geq 1$, let $\psi(i)$ be such that $2^{\psi(i)} \leq \rho \cdot 4^i \leq 2^{\psi(i)+1} - 1$. Note that for any j, $\psi(i) = j$ for at most one value of i. Then, for any $j \geq 1$, $R_j = F_i$ if $j = \psi(i)$ and $R_j = \emptyset$ if there is no i such that $j = \psi(i)$. We set $\alpha = 32/3$; $\pi = \rho \cdot 4^K$; for any $j \geq 1$, $Q_j = (V, W_j)$ where $W_j = \cup_{i-1 \leq r \leq K} 4^{K-r+1} 2^r F_r$ if $R_j \neq \emptyset$ and $j = \psi(i)$, and $W_j = \emptyset$ if $R_j = \emptyset$.

The following lemma proves π-connectivity.

Lemma 6. *With probability at least $1 - 4/n$, every edge $e \in F_i = R_{\psi(i)}$ for each $i \geq 1$ is $\rho \cdot 4^K$-heavy in $Q_{\psi(i)}$.*

Proof. Consider any edge $e \in F_i$. Since $F_i \subseteq Y_{i-1}$, Lemma 3 ensures that e is $\rho \cdot 2^i$-heavy in $G_{i-1} = (V, X_{i-1})$, and therefore $\rho \cdot 2^{2i-1}$-heavy in $(V, 2^{i-1} X_{i-1})$. Since $\varepsilon \leq 1$, Lemma 5 ensures that with probability at least $1 - 4/n$, the weight of each cut in $(V, 2^{i-1} X_{i-1})$ is preserved up to a factor of 2 in $Z_i = (V, \cup_{i-1 \leq r \leq K} 2^r F_r)$. Thus, e is $\rho \cdot 4^{i-1}$-heavy in Z_i.

Consider any cut C containing $e \in F_i$. We need to show that the weight of this cut in $Q_{\psi(i)}$ is at least 4^K. Let the maximum λ_a of an edge a in C be $\rho \cdot 4^{k_C}$, for some $k_C \geq i$. By the above proof, a is $\rho \cdot 4^{k_C - 1}$-heavy in Z_{k_C}. Then, the total weight of edges crossing cut C in $Q_{\psi(k_C)}$ is at least $\rho \cdot 4^{k_C - 1} \cdot 4^{K - k_C + 1} = \rho \cdot 4^K$. Since $k_c \geq i$, $\psi(k_C) \geq \psi(i)$ and $Q_{\psi(k_C)}$ is a subgraph of $Q_{\psi(i)}$. Therefore, the the total weight of edges crossing cut C in $Q_{\psi(i)}$ is at least $\rho \cdot 4^K$. □

We now prove the α-overlap property. For any cut C, let $f_i^{(C)}$ and $w_i^{(C)}$ respectively denote the total weight of edges crossing cut C in F_i and $W_{\psi(i)}$ respectively for any $i \geq 0$. Further, let the number of edges

crossing cut C in $\cup_{i=0}^{K} 2^i F_i$ be $f^{(C)}$. Then,

$$\sum_{i=1}^{K} \frac{w_i^{(C)} 2^{\psi(i)-1}}{\pi} \leq \sum_{i=1}^{K} \frac{w_i^{(C)} \rho \cdot 4^i}{2\rho \cdot 4^K} = \sum_{i=1}^{K} \frac{w_i^{(C)}}{2 \cdot 4^{K-i}} = \sum_{i=1}^{K} \frac{w_i^{(C)}}{2 \cdot 4^{K-i}} = \sum_{i=1}^{K} \sum_{r=i-1}^{K} \frac{f_r^{(C)} \cdot 2^r \cdot 4^{K-r+1}}{2 \cdot 4^{K-i}}$$
$$= \sum_{i=1}^{K} \sum_{r=i-1}^{K} \frac{f_r^{(C)}}{2^{r-2i-1}} = \sum_{r=0}^{K} \sum_{i=1}^{r+1} \frac{f_r^{(C)}}{2^{r-2i-1}} = \sum_{r=0}^{K} \frac{f_r^{(C)}}{2^r} \sum_{i=1}^{r+1} 2^{2i+1} = \frac{32}{3} \sum_{r=0}^{K} 2^r f_r^{(C)} = \frac{32}{3} f^{(C)}.$$

Using Theorem 4, we conclude the proof of Theorem 7.

Size of the skeleton graph. We now prove that the expected number of edges in G_ε is $O(n \log n/\varepsilon^2)$. For $i \geq 1$, define D_i to be the set of connected components in the graph $G_i = (V, X_i)$; let D_0 be the single connected component in G. For any $i \geq 1$, if any connected component in D_i remains intact in D_{i+1}, then there is no edge from that connected component in F_i. On the other hand, if a component in D_i splits into η components in D_{i+1}, then the algorithm explicitly ensures that $\sum_{e \in F_i} \frac{w_e}{\lambda_e}$ from that connected component is $\sum_{e \in F_i} \frac{2^i}{\rho \cdot 4^i} \leq \left(\frac{\rho \cdot 2^{i+2} \cdot 2^i}{\rho \cdot 4^i} \right) \eta = 4\eta \leq 8(\eta - 1)$. Therefore, if $d_i = |D_i|$, then

$$\sum_{i=1}^{K} \sum_{e \in F_i} \frac{w_e}{\lambda_e} \leq \sum_{i=1}^{K} 8(d_{i+1} - d_i) \leq 8n,$$

since we can have at most n singleton components. It follows from Theorem 4 that the expected number of edges added to G_ε by the sampling is $O(n \log n/\varepsilon^2)$. Since the number of edges added to G_ε in steps 6 and 7 of the algorithm is $O(n \log n/\varepsilon^2)$, the total number of edges in G_ε is $O(n \log n/\varepsilon^2)$.

Time complexity of the algorithm. If $m \leq 2\rho n$, the algorithm terminates after the first step which takes $O(m)$ time. Otherwise, we prove that the expected running time of the algorithm is $O(m + n \log n/\varepsilon^2) = O(m)$ since $\rho = \theta(\log n/\varepsilon^2)$. First, observe that phase 1 takes $O(m + n \log n)$ time. We will show that iteration i of phase 2 takes $O(|Y_{i-1}|)$ time. Since $Y_i \subset X_i$ and $\mathbb{E}[|X_i|] = \mathbb{E}[|X_{i-1}|]/2$, and $|Y_0| \leq m$, it follows that the expected overall time complexity of phase 2 is $O(m)$. Finally, the time complexity of phase 3 is $O(m + n \log n/\varepsilon^2)$ (see e.g. [10]).

In iteration i of phase 2, the first step takes $|Y_{i-1}|$ time. We show that all the remaining steps take $O(|X_i| + n \log n)$ time. Since $X_i \subseteq Y_{i-1}$ and the steps are executed only if $Y_{i-1} = \Omega(n \log n/\varepsilon^2)$, it follows that the total time complexity of iteration i of phase 2 is $O(|Y_{i-1}|)$.

First, observe that step 8 and the if block of step 12 take $O(|X_i|)$ time. So, we are left with the repeated invocations of the else block of step 12. Each iteration of the else block takes $O(|V_c| \log n + |E_c|)$ time for the current V_c, E_c. So, the last invocation of the else block takes at most $O(|X_i| + n \log n)$ time. In any other invocation, $|E_c| = \Omega(|V_c| \log n)$ and hence the time spent is $O(|E_c|)$. We show that $|E_c|$ decreases by a factor of 2 from one invocation of the else block to the next; then the total time over all invocations of the else block is $O(|X_i| + n \log n)$.

To see that the $|E_c|$ halves from one invocation of the else block to the next, consider an iteration that begins with $|E_c| > 2k_i \cdot |V_c|$. By Lemma 2, E_c for the next iteration (denoted by E_c') comprises only edges in the first k_i NI forests constructed in the current iteration. So $|E_c'| \leq k_i \cdot |V_c| < |E_c|/2$.

4 Future Work

The obvious open question is whether these results can be extended to weighted graphs, at least if the weights are polynomially bounded in n. Another possibility is to extend these results to the semi-streaming model for unweighted graphs. A more ambitious open problem is to obtain an efficient (i.e. near-linear in m) algorithm that constructs a skeleton containing $o(n \log n)$ edges while approximately preserving the weights of all cuts with high probability.

References

[1] Kook Jin Ahn and Sudipto Guha. Graph sparsification in the semi-streaming model. In *ICALP (2)*, pages 328–338, 2009.

[2] Joshua D. Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-Ramanujan sparsifiers. In *STOC*, pages 255–262, 2009.

[3] András A. Benczúr and David R. Karger. Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time. In *STOC*, pages 47–55, 1996.

[4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.

[5] Wai Shi Fung and Nicholas J. A. Harvey. Graph sparsification by edge-connectivity and random spanning trees. *CoRR*, abs/1005.0265, 2010.

[6] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. Graph sparsification via refinement sampling. *CoRR*, abs/1004.4915, 2010.

[7] Andrew V. Goldberg and Satish Rao. Beyond the flow decomposition barrier. *J. ACM*, 45(5):783–797, 1998.

[8] Navin Goyal, Luis Rademacher, and Santosh Vempala. Expanders via random spanning trees. In *SODA*, pages 576–585, 2009.

[9] Ramesh Hariharan and Debmalya Panigrahi. A general framework for graph sparsification. *CoRR*, abs/1004.4080, 2010.

[10] Voratas Kachitvichyanukul and Bruce W. Schmeiser. Binomial random variate generation. *Commun. ACM*, 31(2):216–222, 1988.

[11] David R. Karger. Random sampling in cut, flow, and network design problems. In *STOC*, pages 648–657, 1994.

[12] David R. Karger. Using randomized sparsification to approximate minimum cuts. In *SODA*, pages 424–432, 1994.

[13] Rohit Khandekar, Satish Rao, and Umesh V. Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4), 2009.

[14] Tom Leighton and Ankur Moitra. Extensions and limits to vertex sparsification. In *STOC*, page To appear, 2010.

[15] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *FOCS*, pages 3–12, 2009.

[16] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1997.

[17] Hiroshi Nagamochi and Toshihide Ibaraki. Computing edge-connectivity in multigraphs and capacitated graphs. *SIAM J. Discrete Math.*, 5(1):54–66, 1992.

[18] Hiroshi Nagamochi and Toshihide Ibaraki. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica*, 7(5&6):583–596, 1992.

[19] Jonah Sherman. Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$-approximations to sparsest cut. In *FOCS*, pages 363–372, 2009.

[20] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. In *STOC*, pages 563–568, 2008.

[21] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *STOC*, pages 81–90, 2004.

[22] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *CoRR*, abs/cs/0607105, 2006.

A Proof of Lemma 1

To prove Lemma 1, we will need two theorems from [9]. The first theorem is a non-uniform extension of Chernoff bounds. [2]

Theorem 8 (Hariharan-Panigrahi[9])**.** *Consider any subset C of unweighted edges, where each edge $e \in C$ is sampled independently with probability p_e for some $p_e \in [0,1]$ and given weight $1/p_e$ if selected in the sample. Let the random variable X_e denote the weight of edge e in the sample; if e is not selected in the sample, then $X_e = 0$. Then, for any p such that $p \le p_e$ for all edges e, any $\varepsilon \in (0,1]$, and any $N \ge |C|$, the following bound holds:*[3]

$$\mathbb{P}\left[|\sum_i X_e - |C|| > \varepsilon N\right] < 2e^{-0.38\varepsilon^2 pN}.$$

To state the second theorem, we need the following definitions.

Definition 3. *For any undirected graph G and for any $k > 0$, the k-projection of any cut C is the set of k-heavy edges in C.*

Definition 4. *The* edge connectivity *of an undirected graph G is the minimum weight of a cut in G.*

The theorem counts the number of distinct k-projections in cuts of weight αk for any $k \ge c$, where c is the edge connectivity of the graph.

[2] For Chernoff bounds, see e.g. [16].

[3] For any event $\mathscr{E}$, $\mathbb{P}[\mathscr{E}]$ represents the probability of event $\mathscr{E}$.

Theorem 9 (Hariharan-Panigrahi[9])**.** *For any undirected graph with edge connectivity c and for any $k \geq c$ and any $\alpha \geq 1$, the number of distinct k-projections of cuts of weight at most αk is at most $n^{2\alpha}$.*

Using the above two theorems, we now prove Lemma 1.

Proof of Lemma 1. Let $\mathscr{C}_j$ be the set of all cuts C such that $2^j\pi \leq r^{(C)} \leq 2^{j+1}\pi - 1,\ j \geq 0$. We will prove that with probability at least $1 - 2n^{-2^{j+1}}$, all cuts in $\mathscr{C}_j$ satisfy the property of the lemma. Then, the lemma follows by using the union bound over j since $2n^{-2} + 2n^{-4} + \ldots + 2n^{-2j} + \ldots \leq 4n^{-2}$.

We now prove the property of the lemma for cuts $C \in \mathscr{C}_j$. Since each edge $e \in R^{(C)}$ is sampled with probability p in obtaining $\widehat{R^{(C)}}$, we can use Theorem 8 with sampling probability p. Then, for any $R^{(C)}$ where $C \in \mathscr{C}_j$, by Theorem 8, we have

$$\mathbb{P}\left[\left|\widehat{r^{(C)}} - r^{(C)}\right| > \delta q^{(C)}\right] < 2e^{-0.38\cdot\delta^2\cdot p\cdot q^{(C)}} \leq 2e^{-0.38\cdot\delta^2\cdot p\cdot\pi\cdot 2^j} \leq 2e^{-6\cdot 2^j \ln n} = 2n^{-6\cdot 2^j},$$

since $q^{(C)} \geq \pi \cdot 2^j$ for any $C \in \mathscr{C}_j$. Since each edge in $R^{(C)}$ is π-heavy in (V,Q), Theorem 9 ensures that the number of distinct $R^{(C)}$ sets for cuts $C \in \mathscr{C}_j$ is at most $n^{2\left(\frac{\pi\cdot 2^{j+1}}{\pi}\right)} = n^{4\cdot 2^j}$. Using the union bound over these distinct $R^{(C)}$ edge sets, we conclude that with probability at least $1 - 2n^{-2^{j+1}}$, all cuts in $\mathscr{C}_j$ satisfy the property of the lemma. $\square$